Variational autoencoder

Autoencoders - Issues

- No sampling
- Learns just compression

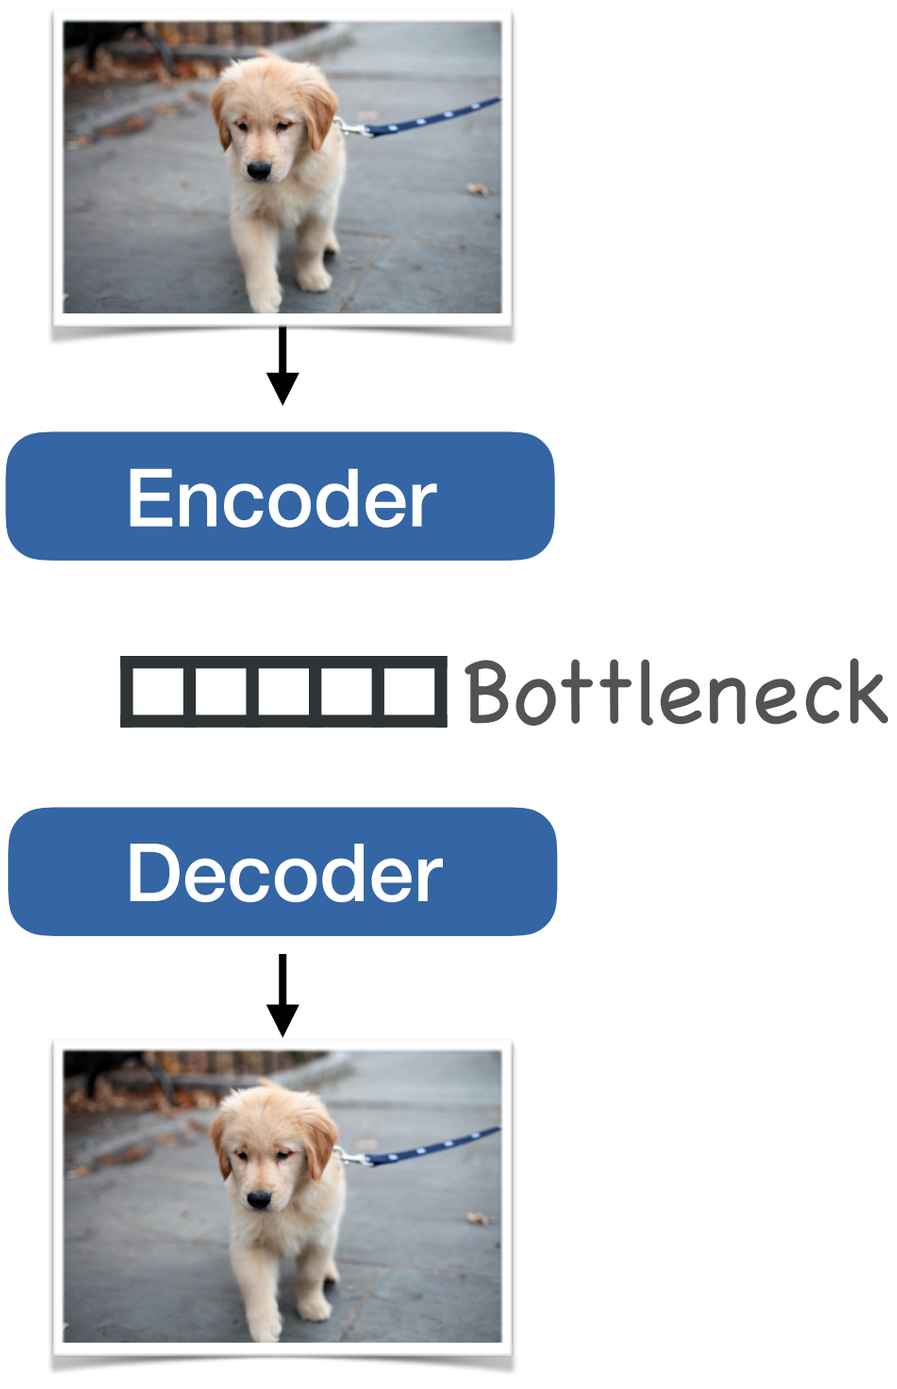

Variational autoencoder

- Autoencoder
  - with noise in bottleneck

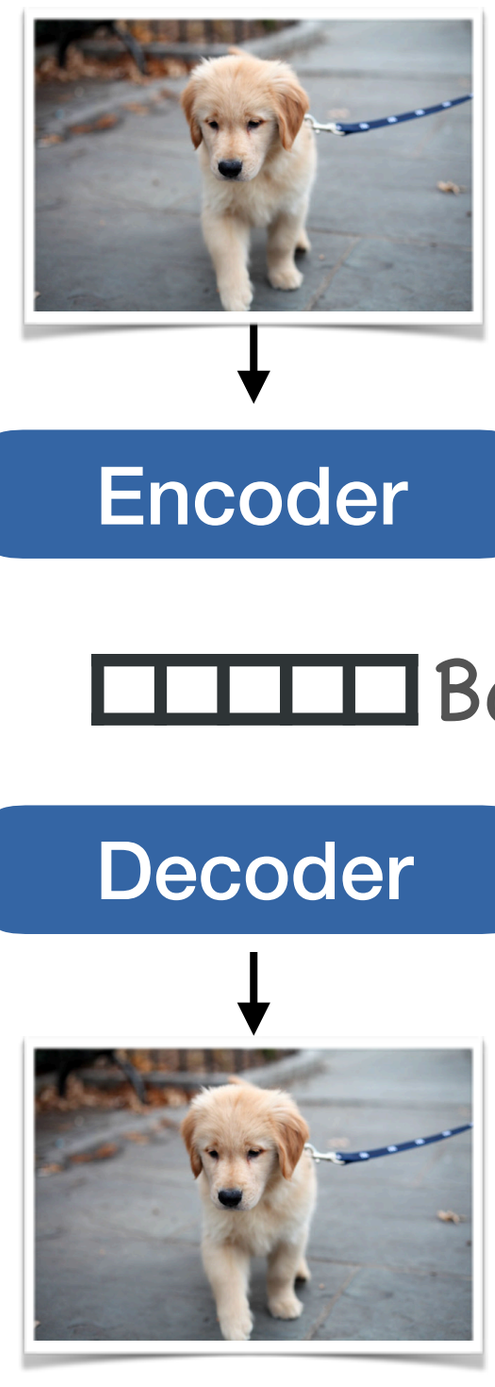

Why does noise help?

Variational autoencoder - formal definition

- Encoder
  - $q(\mathbf{z}\,|\,\mathbf{x}) = \mathcal{N}\left(\mathbf{z}; \mu_\theta(\mathbf{x}), \sigma^2_\theta(\mathbf{x})\mathbf{I}\right)$
- Sampling $\mathbf{f} \sim q(\mathbf{z}\,|\,\mathbf{x})$
- Decoder
  - $P(\mathbf{x}\,|\,\mathbf{f})$
- Approximately learns $P(\mathbf{x})$
  - Variational lower bound

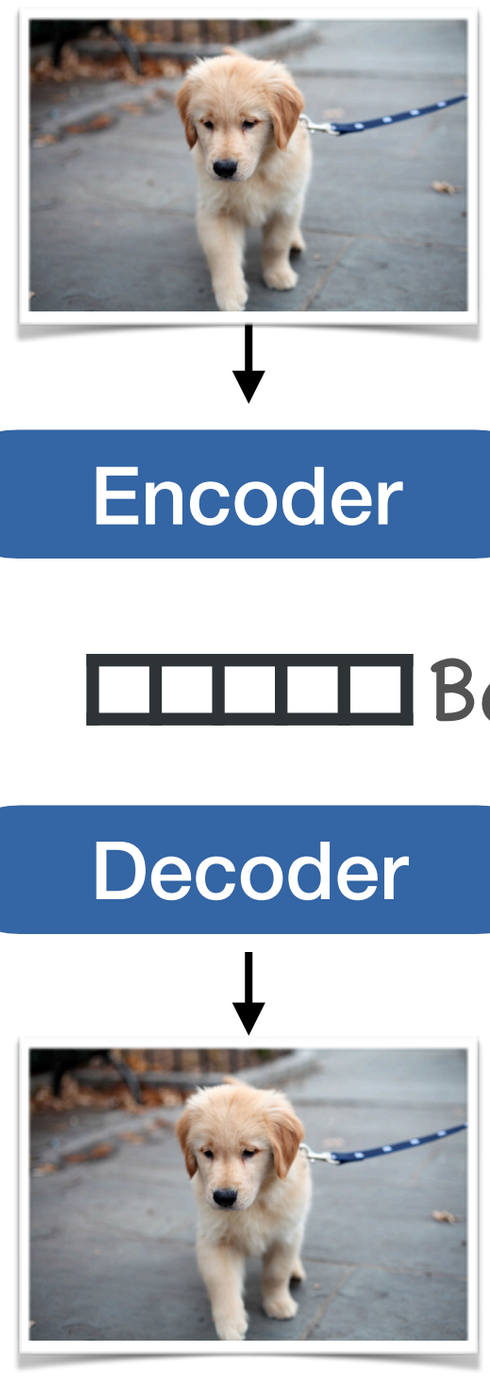

Variational autoencoder - Issues

- Fails in high dimensions
  - Hard to embed spherical distributions
- Blurry outputs
  - Pixel-distance